BUSINESS ISSUES IN TECHNOLOGY
ComputerLetter
Published by Technologic Partners

Volume 10, Number 4
January 31, 1994

The Video Game

On-demand video networks provide a billion-dollar opportunity for computer makers able to meet some stringent requirements

Sometimes we marvel at the bravado that allows start-up companies to launch a new product with little or no evidence anybody will pay good money for it. But the fledgling entrepreneurs infesting Silicon Valley have nothing on the cable and telephone giants when it comes to audacity. During the past year, several service providers have conducted focus groups and market trials to gauge consumer appetites for home shopping, video on demand, network gaming, and other services imagined for the digital information superhighway. Despite apparently inconclusive results from these trials, providing limited service to dozens of cooperative participants, many are nevertheless forging ahead with ambitious plans to deploy largely untested technologies to thousands of paying customers.

In the latest announcement along those lines, **Pacific Telesis** has said it will roll out an on-demand video service that will pass — or reach — 750,000 homes in California by the end of next year. It's too early to tell how receptive the public will be to all these new services, or even exactly what form the new services will take. But we can take for granted the determination of the carriers to rewire the nation's communication infrastructure, allowing us to explore the role to be played in this effort by the computer industry.

A few months back we looked at the prospects for the advanced set-top boxes subscribers will need to travel the video highway ("All Aboard the SST," Sept. 27, 1993) and concluded that they might represent a $1 billion market by 1997. We plan to examine the video opportunity for relational and object-oriented database software in a future letter. This week's sermon, again prepared with the help of our friends at Smaby Group, concerns video servers — sophisticated computer platforms that will pump video signals to subscribers, store and manage databases, monitor traffic, track usage, verify credit, and facilitate billing.

Counting on another billion-dollar opportunity, virtually all of the big names in big iron have stated their intention to capture a fair share of this important new market. The potential is indeed large. Video server revenues should start out at about $50 million this year, jump to $375 million in 1995, and hit the $1 billion mark in 1998. That much is evident from an analysis of the plans now being discussed and some reasonable assumptions about expansion. What is not so clear is the nature of the equipment that will predominate and who will provide it.

While there's plenty of talk about convergence between cable television and telephony, the biggest customers for large-scale video server equipment initially are likely to be the regional telephone companies. It's a matter of geography and economics. The telephone companies typically cover huge metropolitan areas or even entire states. Cable franchises are usually limited by the borders of a municipality; the adjacent town

(Continued on Page Two)

This Week

Private Profiles:

VIDEO SERVERS

Continued from Page One

is likely to be served by a different cable system owned by a different operator. Further, the cable operators don't enjoy the financial resources that the telephone companies have, although consolidations such as the proposed **Bell Atlantic** acquisition of **Tele-Communications** could change that picture. So, while they operate in a less restrictive regulatory environment, cable companies still face significant constraints. Even on a local level, however, they are likely to use the same kind of server technology to upgrade their networks.

Fiber to the people

Hybrid networks of optical fiber and coaxial cable will carry video streams to the home.

A number of local telephone companies, including most of the Baby Bells, have announced plans to offer video dial tone service — the transport mechanism for delivering switched, interactive video connections — in selected areas. There are two basic mechanisms being considered for delivery of video dial tone: Asynchronous Digital Subscriber Line (ADSL), a Bellcore technology that sends a compressed digital video signal over an existing copper phone line; and a hybrid optical fiber/coaxial cable system. The ADSL approach would allow the phone companies to leverage their installed wiring plant, but the jury is still out on how cost-effective the electronics will be for getting high-speed data on and off the wire.

Stream power

Based on FCC filings to date, the hybrid fiber and coax architecture is more popular but will require a heavy investment in fiber lines that could slow down the deployment of video dial tone. This scheme uses newly installed fiber-optic cables to carry video streams from a central office to so-called optical network units located in a neighborhood or serving area; at that point, signals are converted from optical to electrical pulses and switched onto coaxial cable for delivery to individual homes.

The computer industry's opportunity lies with products that will supply this plumbing system, known in FCC parlance as a Level 1 network, with streams of video information. As the FCC defines things, it is the Level 2 providers — there

BUSINESS ISSUES IN TECHNOLOGY
ComputerLetter
A Technologic Publication

Editor: Richard A. Shaffer
419 Park Avenue South • Suite 500
New York, NY 10016
Tel: 212 696-9330
Fax: 212 696-9793 • MCI: 371-2412

Managing Editor: John W. Wilson
P.O. Box 869
Larkspur, CA 94977-0869
Tel: 415 924-1274
Fax: 415 924-0945 • MCI: 381-8419

Contributors: Paulina Borsook, Susan Carlson, Brian O'Connell, Tom Sato, Gary Smaby, Kathleen K. Wiegner

Art Director: Natalie Tilton
Production Assistant: Thaddeus Batt
Circulation Director: John Lo

Subscriptions are $595 per year in the United States. Outside the United States subscriptions are $695. For subscription inquiries, or to order back issues and reprints, contact John Lo at Technologic in New York.

Technologic Partners

could be several in a given area — that will develop and supply information services over the network. Level 2 service providers will have to communicate control information to and from the Level 1 network, perhaps via an Ethernet-type interface, as well as deliver compressed video information, encoded with its intended destination, over one or more high-speed connections.

The video server and the set-top devices used to request, tune, and decode the signals are all the responsibility of the Level 2 service provider. Initially, at least, many of the Level 1 providers will also handle Level 2 services. Several of the Baby Bells have created information services subsidiaries that will buy network services from the sister company and perform the service development and system integration tasks necessary to deploy a complete system. Given that few standards have yet emerged, the eventual arrival of additional Level 2 providers will create some confusion and, no doubt, some irritation for consumers as they discover that an additional set-top device is required each time they sign up with a new service. As far as the video server is concerned, though, it appears that most systems will use compression schemes based on the pending MPEG-2 standard.

Conventional computer systems will strain under the demands imposed by video on demand.

Video librarian

The basic function of the video server is to retrieve compressed video data from an online library and route it to a requesting subscriber. The MPEG-2 proposal is flexible regarding the data rate for a single video stream; for this discussion, we assume that 4 megabits per second will be required for satisfactory picture quality. That means a minute of video, stored in compressed form, will occupy 30 megabytes on disk; a 100-minute movie, 3 gigabytes. Storing a few hundred movies, plus educational videos and clips associated with, for example, shopping functions, could easily drive storage needs toward a terabyte and beyond. Obviously, this requirement won't be met by constructing enormous hard disk farms. Rather, we expect to see three-tier systems in which the dozen or so titles in greatest demand are stored in random access memory, perhaps another 50 to 100 on disks, and hundreds more on tape.

Streaming that data out to subscribers is also a demanding task. If the video server is to generate 1,000 video streams (imagine each of 1,000 customers viewing a unique spot in a movie selection), that means formatting and outputting data to the Level 1 network at a rate of 4 gigabits per second. And, unless the video data is cached in memory, the same amount of data will have to be read from disk each second. Thus, a video server handling 1,000 simultaneous streams would have to have an internal bandwidth (between disk, memory, and I/O subsystem) capable of supporting the movement of nearly a gigabyte of data per second, plus any necessary overhead.

Grunts and groans

It seems obvious that conventional computer systems will strain under the demands imposed by even a modest-sized video-on-demand network. The most powerful RISC-based systems, such as those proposed by **IBM**, **Silicon Graphics**, and **Sun Microsystems** for this application, will have to be used in clusters to serve anything bigger than a small village. Supporting a thousand simultaneous video streams would be difficult even for Silicon Graphics' Challenge server, which boasts an internal bus speed of 1.2 gigabytes per second. Several of these systems will be tied together to meet the needs of **Time-Warner**'s Orlando, Fla., trial.

IBM has offered its RS/6000 systems, including the POWERParallel cluster line, as candidates for video duty. But at a December trade show for the cable industry IBM was showing its ES/9000 mainframe in that role. By distributing various tasks to intelligent subsystems — letting an input/output controller handle the packetizing and synchronization required to get data onto the optical network, for example — IBM believes it can configure its mainframes to meet some of the most demanding requirements.

How to beat the VCR

It may be that entirely new system architectures will be required to move the massive amounts of data video servers will handle. That is the contention of **Hewlett-Packard**, which will supply Pacific

Screen Test

As subscriber penetration grows...

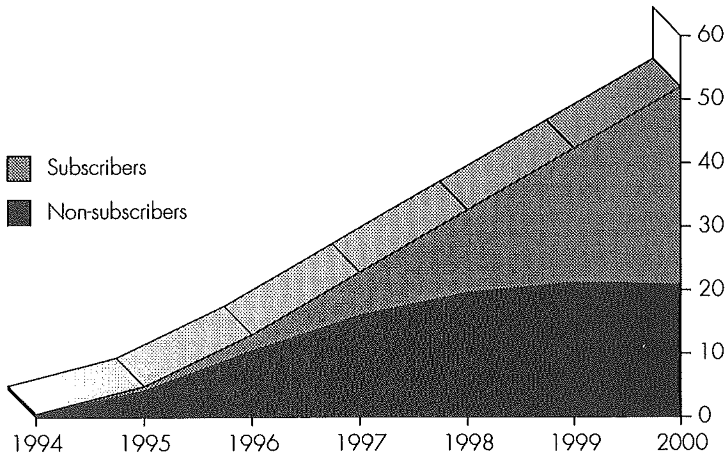

Total number of homes passed by video-on-demand service in millions. Includes subscribers (gray area) and non-subscribers (black area).

...demand for video equipment...

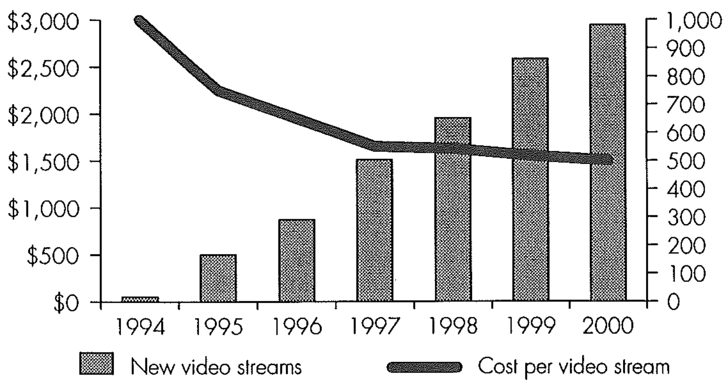

Average cost per video stream in dollars (line, left scale); shipments of video servers as measured in thousands of video streams handled (bars, right scale).

...will build a significant market

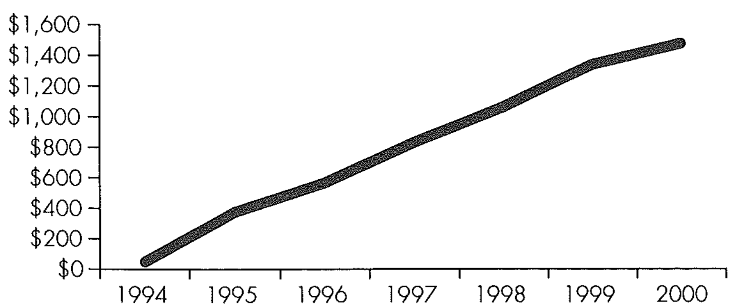

Sales of video servers, including data storage equipment, in millions of dollars.

Data for all charts: Smaby Group forecasts

Telesis with servers for its video-on-demand network, starting late this year. HP has created a Video Communications Division to develop a new I/O subsystem it calls a video transfer engine, which is expected to achieve transfer rates up to 3 gigabits per second. This will enable the server to support fast-forward or rewind at speeds as high as 10 times the standard movie speed — an important requirement if video on demand is to compete with the VCR showing rented tapes.

Makers of massively parallel processing systems, including **Kendall Square Research**, **MasPar**, and **nCUBE**, see the video server as a perfect fit with their design approach. As we've reported, a spin-off of the David Sarnoff Research Center, called **Sarnoff Real Time**, is planning to pursue the video server opportunity with MPP technology designed specifically for video and image processing. MPP systems are highly scalable — meaning service providers can start with small systems to handle initial demand and add incremental performance to meet increasing performance. There are at least two ways to use the hundreds of microprocessors in an MPP computer. In one configuration, each processor with its attached disks might be responsible for a particular subset of the movie library; subscribers would be assigned to processors based on their selections. Conversely, each processor might be dedicated to a subset of subscribers, reaching through the internal network of the system to grab the necessary data.

Guess again

The size of the market for video servers will depend on variables that can't be predicted with much confidence, including the rate at which customers will sign up for the service when video dial tone reaches their street, the use they will put to the system, and the efficiency with which the server handles video stream requests. Still, we think Smaby Group has made some reasonable guesses, which we've shown in the three charts on this page.

The first chart assumes, based on statements made by service providers, that about 500,000 homes will be enabled for video on demand by the end of this year — meaning that the service will pass that

Hardware needs will depend in part on how much subscribers make use of the system.

many homes — and that the number will approach 5 million in 1995. After that, the growth rate will depend in part on results of the initial rollout; the assumption behind this forecast is that service providers will generate sufficient revenue to continue the deployment until 52 million homes have been reached by 2000. That would be 80% of the homes now wired for cable TV. Meanwhile, the number of subscribers will start at a modest 50,000 this year, or 10% of the homes passed, and rise to 31 million by the end of the decade, which would represent a 60% penetration.

How much is enough?

The question then becomes: How heavily will subscribers use the system? The model used by Pacific Telesis and HP assumes that 10,000 video streams will suffice to handle the needs of the 100,000 subscribers expected by the end of next year. We think that may be too conservative because the early adopters who will constitute the bulk of the first subscribers are likely to be avid channel surfers. In the end, the service provider will have to determine what the peak load will be and build an infrastructure to support it. The second chart assumes that each video stream will support only three subscribers initially, with the number rising to nine by the end of the decade, producing a requirement for new systems handling 16,000 video streams this year and 980,000 by 2000. Over the same time span, the cost of server equipment, including data storage systems, is likely to fall from an average of $3,000 per stream this year to $1,500. All of which leads to the forecast in the final chart of a hardware market that grows at a rate of 30% a year, reaching $1.5 billion at the end of the decade.

Video on a budget

These will be costly systems, if the estimate of $3,000 per video stream is at all realistic. It would be possible to cut the memory and video stream requirement for a video server, and hence its cost, by offering something less than true video on demand. Some providers are betting that it will be possible to satisfy most couch potatoes — perhaps 80% — with a relatively small selection of new releases and classics. If they're right, so-called near video on demand may make more sense.

In this approach, the same movie is started at intervals of, say, five minutes on each of 20 channels. When a customer requests that movie, his set-top box automatically tunes to the channel offering the next showing and counts down to the start; the average wait will be two and a half minutes. The movie can be "paused" by jumping to the next channel, which is five minutes behind. Under this scheme, ten movies could be supplied to thousands of subscribers with only 200 simultaneous video streams. This uses much less online storage. Routing and switching of individual streams to homes wouldn't be required.

Good enough isn't

But the information superhighway zealots won't hear of such compromises. They believe there's money to be made from time shifting — letting consumers watch *60 Minutes* without having to rearrange their Sunday schedule or fiddle with the VCR, for example. Educational and training videos are another narrowcasting opportunity. And then there are the truly interactive experiences that demand individualized video capability — strolling through a virtual shopping mall, calling up video clips of merchandise; or planning your next vacation, complete with a walk-through of the hotel and a helicopter tour of the beach. If one major provider offers these services successfully, competition will force the rest to follow suit.

For computer vendors, the video server opportunity will likely play out like so many other promising niches before it. Scalability, reliability, software support, and standards compliance will be critical. But service providers will be most concerned about price, especially when they get past the trial stage and start to deploy systems in volume. The hardware side of the video server business will become very cost-competitive after the high-visibility trials have concluded. But that will help video server technology move from Hollywood glitz to Main Street mundane. Specialized (and lower-cost) video or multimedia servers for real-time audio/video applications in corporate communications, education, and training seem a likely extension of this technology. □

Private Profiles: Tetherless Access

Building a wireless on-ramp to the information highway

As long ago as 1972, Dewayne Hendricks and thousands of other college-aged hackers were spending their evenings in computer labs, glued to the ARPAnet (later renamed the Internet). They became the first generation of Internet junkies — Internauts, as they called themselves — with an intricate understanding of gophers, MultiUser Dungeons, and other mysteries of the information superhighway. (For the uninitiated, gophers are programs that tunnel through the Internet for information and MUDs are artificial worlds where Internauts meet, play games, and even get married.) Now, as Mr. Hendricks and his small company, **Tetherless Access**, prepare to install their first community wireless network, that 22-year apprenticeship on the Internet looks like a prime asset. In a wireless communications industry increasingly dominated by giants, Mr. Hendricks maintains that Tetherless is the only company that really understands the Internet culture, its technical intricacies, and the pivotal role it will play in the mobile communications industry.

Tetherless hopes its understanding of Internet culture and technology will set it apart in the wireless wars.

Virtually a company

That's quite a claim for a fledgling company with no wired infrastructure, no financial resources to speak of, and not even a corporate headquarters (the 10 Tetherless employees operate a virtual office linked by communications systems). What Tetherless does have — and the reason we think it's of interest — is a unique approach to Internet access that makes use of wireless communications to provide individuals and communities with direct, high-speed connections at a reasonable cost. We think there may be a lot of potential customers for a wireless service that is based on the TCP/IP protocol suite, the standard language for using the Internet; moves data at 64 kilobits per second (Kbps); and requires an investment of $3,000 or less per site. If Tetherless turns out not to be up to the challenge of delivering all this, others will be.

Outfitting explorers

Of course, such a service is not for everyone. Most people use the Internet today primarily to send and receive electronic mail, and for them a standard dial-up service, or one of the wireless wide area networks — **Ardis**, **RAM Mobile Data**, or **McCaw Cellular Communications'** emerging Airdata — is perfectly adequate. But for those who want to explore the Internet fully and use the software and tools available on it, the speeds of 600 to 2,400 bits per second available on the wireless WANs or even the 14.4 kilobits per second of a wired modem aren't fast enough. Only Airdata implements the common Internet language. A fourth wireless network provider, **Metricom**, claims that its forthcoming UbiquiNet system will operate at 20 to 30 Kbps and possibly faster. For the dedicated Internaut, the 56 Kbps throughput of a high-speed leased line is generally recognized as the minimum acceptable speed.

Until Tetherless came along, the only way to reach that speed was to lease a 56 Kbps dataline from a regional telephone company at a cost of perhaps $4,500 for installation and equipment and about $2,200 per year in leasing charges. An

At a Glance:

TETHERLESS ACCESS

Location:	43730 Vista Del Mar Fremont, CA 94539
Telephone:	(510) 659-0809
CEO:	Dewayne Hendricks
Business:	Wireless access to the Internet
Financing:	$200,000
Investors:	Company principals
Employees:	10
Estimated Sales:	$1.5 million in 1994

alternative, where it is available, would be an Integrated Services Digital Network line, which would run about $200 for installation and $30 per month plus a penny-per-minute usage fee — perhaps $430 per month of round-the-clock use.

Pick the right fight

We wouldn't want to suggest that Tetherless can compete against wired services forever on speed alone. Eventually, high-speed datalines will be widely available at much lower cost and even homes will have easy access to the Internet as telephone companies and cable television operators start to compete in earnest. In the long run, Tetherless' real hope probably lies in its potential in the mobile computing market, where the speed and robust implementation of TCP/IP will set its Net/Access system apart.

Telluride's system will test the use of wireless technology in rural communities.

Another important component of Mr. Hendricks' strategy is to avoid significant capital investments in infrastructure. Taking its model from the egalitarian structure of the Internet, which evolved through private investments by thousands of universities, government laboratories, and high-tech companies, Tetherless believes users of the network will want to own their own share of the infrastructure. So Tetherless' income will be generated by selling its radio systems and software ($3,000 for the package), and by providing services over the network, beginning with direct access to the Internet. Initially, the company will also be responsible for installing the **Sun Microsystems** workstations and **cisco Systems** routers that will link each community network through a T-1 line to the wired Internet — but it doesn't plan to stay in the gateway business. There will be no metered usage charges, but Tetherless will collect an annual $350 support fee as well as monthly Internet access fees.

Wireless country

The Net/Access system seems particularly appropriate for rural and hard-to-reach communities. The first full-scale test of this thesis started this month, when Tetherless began installing a 6-node community network in Telluride, Colo., a community of 1,500 people located 9,000 feet above sea level and surrounded on all sides by 13,000-foot peaks. When it came to installing telephone equipment, Telluride lagged years behind the rest of the country. Until three years ago, most residents were still on party lines. Today, the Telluride InfoZone program is perhaps the country's most progressive community program dedicated to deploying new communications technologies.

With sizable grants from government and industry (including **Apple Computer** and **IBM**), Telluride is installing a network of public-access computer systems in its schools, libraries, museum, town buildings, and medical facility all linked via the wireless Net/Access system. The Telluride network is actually the fourth that Tetherless has installed. In conjunction with Apple, it has deployed pilot networks in San Diego (6 nodes), the San Francisco Bay Area (15 nodes in the Smart Valley project), and in Washington, D.C. (2 nodes). It has also been chosen to install a 50-node community network in Davis, Calif. So far, the company is providing service free to these communities, although it plans to begin offering commercial service in most of the areas soon.

Telluride, long neglected by its regional phone company, favors the community investment philosophy behind Tetherless. That approach will allow communities like Telluride to reinvest locally in order to expand their own communications infrastructure, rather than paying a regional phone company "to take our money and invest it in something like trans-Asiatic fiber systems," says Richard Lowenberg, director of Telluride's InfoZone Program. The community is also ideally suited to the idea of wireless Internet access. Telluride has a highly educated population with more computers per capita than any other non-university or non-corporate town.

What a ham

The Net/Access system is based on code that Mr. Hendricks, a licensed amateur radio operator from the age of 12, began writing in 1986 to connect his Macintosh into the Internet using an amateur radio modem. The program was crafted to meet the software requirements of the Internet and all the technical demands of the noisy,

unlicensed ISM (industrial, scientific, and medical) radio band. To know whether the software would work, it had to be deployed on a wide scale. So Mr. Hendricks posted his program for free on the Internet, eventually attracting users from around the world who were willing to invest in radios to test the software.

Since then, the software has evolved considerably. It now totals 170,000 lines of code that connect radio-modems in a fully distributed wireless network, each having direct access to the Internet and to each other. Technically, the network uses spread-spectrum packet-switching radio technology in a peer-to-peer, fully-interconnected, mesh network architecture. Like the Internet itself, there is no central intelligence in the network. Instead, software for routing and managing the network is distributed across the entire grid of network nodes. If any node were to crash or to be removed, the network would simply reconfigure itself around the missing node. That capability, called dynamic routing, is a key feature of the Internet.

Net/Access uses hybrid signaling and other techniques to move data at impressive speeds.

Want to drag?

The speed of the Net/Access system is impressive — if it checks out in the field. Mr. Hendricks claims that analyses conducted in laboratory simulations have demonstrated that the network can sustain a throughput of 64 Kbps, or twice as fast as Metricom's UbiquiNet, promised for later in the first half of the year, which is very similar in design. We've heard plenty of throughput claims that didn't stand up to testing by customers. But Mr. Hendricks cites several technical advantages for Net/Access:

- Using a hybrid radio-signaling technique that combines two distinct spread-spectrum technologies: direct-sequence, in which a single message is chopped into pieces and spread over a range of frequencies; and frequency-hopping, in which the pieces jump from frequency to frequency at regular intervals. The result is supposed to be a better throughput rate than Metricom can achieve by using only frequency-hopping techniques.

- Providing better resistance to signal interference by investing in more sophisticated radio transceivers than Metricom uses. However, at $3,000 the transceivers are far more expensive than Metricom's $700 units.

- Operating in two ISM bands: the 902-928 KHz band — the most heavily-used of the ISM bands, and the only band used by Metricom — and in the less-congested 2.4 MHz band, where higher data rates can be achieved.

Helping hands

Tetherless hopes to reduce the price of its high-speed transceiver quickly with help from **Cylink**, a Sunnyvale, Calif.-based developer of spread-spectrum equipment that makes the AirLink MP radio, which is used in current Net/Access installations. As part of an alliance with Tetherless, Cylink is designing a new family of radio modems for the Net/Access system that Mr. Hendricks hopes will range in price from $3,000 to under $1,000. Cylink, meanwhile, will gain the right to sell the product line in its own market, the bypass services that provide wireless alternatives to telephone company landlines.

Money talks

What Tetherless needs most is financing. Self-funded for the past three years by Mr. Hendricks and other principals, the company is talking with several private individuals and venture capitalists who are interested in its technology. Meanwhile, Tetherless is trying to get by with alliance relationships. Aside from Cylink, Mr. Hendricks is working with **Pandora Systems** for interface technology connecting commercial software to the Internet and with **InterNex Information Systems**, which is providing wired access points to the Internet.

For Hendricks and his partners, thinking big seems to be almost a company requirement. We wouldn't be surprised to find them at the FCC's May auction of dedicated radio spectrum for Personal Communications Systems. A PCS license would be a valuable asset in the wireless wars. And after all, as a small company steeped in community network technology, Tetherless would not be required to pay cash. □